arXiv:quant-ph/0512085v1 12 Dec 2005

Random measurement bases, quantum state distinction and applications to the hidden subgroup problem

Pranab Sen
NEC Laboratories America, Inc.
4 Independence Way, Suite 200,
Princeton, NJ 08540, U.S.A.
pranab@nec-labs.com

Abstract

We show that measuring any two quantum states by a random POVM, under a suitable definition of randomness, gives probability distributions having total variation distance at least a universal constant times the Frobenius distance between the two states, with high probability. In fact, if the Frobenius distance between the two states is not too small and their ranks are not too large, even a random orthonormal basis works as above. Since a random POVM is independent of the two states, the above result gives us the first sufficient condition and an information-theoretic solution for the following quantum *state distinction problem*: given an a priori known ensemble of quantum states, is there a single measurement basis, or more generally a POVM, that gives reasonably large total variation distance between every pair of states from the ensemble? Large pairwise trace distance is a trivial necessary condition for the existence of a single distinguishing measurement for an ensemble; however, it is not sufficient, as seen for example by the recent work of Moore, Russell and Schulman [MRS05] on hidden subgroups of the symmetric group. Our random POVM method gives us the first information-theoretic upper bound on the number of copies required to solve the quantum *state identification problem* for general ensembles, i. e., given some number of independent copies of a quantum state from an a priori known ensemble, identify the state. Moreover, this upper bound is achieved by a *single register* algorithm, i. e., the algorithm measures one copy of the state at a time, followed by a classical post-processing on the observed outcomes in order to identify the state.

The standard quantum approach to solving the hidden subgroup problem (HSP), which includes Shor's algorithms for factoring and discrete logarithm, is a special case of the state identification problem where the ensemble consists of so-called *coset states* of candidate hidden subgroups. Combining Fourier sampling with our random POVM result gives us single register algorithms using polynomially many copies of the coset state that identify hidden subgroups having polynomially bounded rank in every representation of the ambient group. In particular, we get such single register algorithms when the hidden subgroup forms a Gel'fand pair, e.g. dihedral, affine and Heisenberg groups, with the ambient group, i. e., the rank in every representation is either zero or one. These HSP algorithms complement earlier results about the powerlessness of random Fourier sampling when the ranks are exponentially large, which happens for example in the HSP over the symmetric group. The drawback of random Fourier sampling based algorithms is that they are not efficient because measuring in a random basis is not. This leads us to the open question of efficiently implementable pseudo-random measurement bases.

1 Introduction

The hidden subgroup problem (HSP) is a central problem in quantum algorithms. Many important problems like factoring, discrete logarithm and graph isomorphism reduce to special cases of the HSP. Almost all

exponential speedups that have been achieved in quantum computing are obtained by solving some instances of the HSP. The HSP is defined as follows: Given a function $f : G \to S$ from a group G to a set S that is constant on left cosets of some subgroup $H \leq G$ and distinct on different cosets, find a set of generators for H. Ideally, we would like to find H in time polynomial in the input size, i. e. $\log |G|$. Almost all efficient quantum algorithms for solving special cases of the HSP, including Shor's algorithms for factoring and discrete logarithm [Sho97], use the same generic approach sometimes called the *standard method*. The standard method for the HSP can be described as follows: evaluate the function f in superposition and ignore the function value to get a state of the form $\sigma_H := \frac{1}{|G|} \sum_{g \in G} |gH\rangle\langle gH|$, where $|gH\rangle := \frac{1}{\sqrt{|H|}} \sum_{h \in H} |gh\rangle$, i .e., σ_H is a uniform mixture of uniform superpositions over left cosets gH of the hidden subgroup H. A state of the form σ_H for some subgroup $H \leq G$ is called a *coset state*. The above procedure can be repeated t times to get t independent copies of the state σ_H. The aim now is to identify H from $\sigma_H^{\otimes t}$.

The coset state based approach to the HSP leads us to consider the following general problem called *quantum state identification*. Given $\sigma_i^{\otimes t}$ from an a priori known ensemble $\mathcal{E} = \{\sigma_1, \ldots, \sigma_m\}$ of quantum states in $\mathbb{C}^n$, identify i. A related problem is the following *quantum state distinction* problem: is there a single measurement basis or more generally a POVM $\mathcal{M}$, that gives reasonably large total variation distance between every pair of states in $\mathcal{E}$? The important point here is that we want a single measurement $\mathcal{M}$ that works well for every pair of states. A solution to the state identification problem trivially gives a solution to the state distinction problem. It is not hard to see that the converse is also true: a POVM $\mathcal{M}$ with *distinguishing power* δ, i.ė., $\mathcal{M}$ solves the state distinction problem with total variation distance at least δ between every pair of states from $\mathcal{E}$, gives an algorithm that identifies the given state with constant probability from $t = O\left(\frac{\log m}{\delta^2}\right)$ independent copies. This algorithm is in fact a *single register* algorithm in that it applies t independent copies of $\mathcal{M}$ to the given $\sigma_i^{\otimes t}$ and does a classical 'minimum-finding style' post-processing on the observed outcomes to guess i. Single register algorithms may have advantages over multi-register algorithms in the interests of efficiency and ease of design; observe that the complexity of a generic k-register measurement increases exponentially with k.

In this work, we study information-theoretic aspects of the general state distinction problem, and use it as a tool for solving the corresponding state identification problem. We also analyse various implications of these two problems, including consequences for the HSP. Our main objective is to find sufficient conditions on the ensemble $\mathcal{E}$ to guarantee the existence of a measurement with distinguishing power δ. It is known that two quantum states can be δ-distinguished by a measurement if and only if they have trace distance at least δ. In general, this measurement depends upon the pair of states to be distinguished. Thus, this result does not give us any way to come up with a single measurement $\mathcal{M}$ is that works well for every pair of states. However, it does provide a necessary condition: in order for a POVM with distinguishing power δ to exist, every pair of states in $\mathcal{E}$ must have trace distance at least δ. On a concrete note, we show that the ensemble of coset states for subgroups of a group G indeed has minimum pairwise trace distance of 1. However, constant pairwise trace distance is not sufficient for the existence of a polynomially distinguishing measurement, as seen for example by the recent work of Moore, Russell and Schulman [MRS05] on hidden subgroups of the symmetric group.

Random POVM and Frobenius distance: In this paper, we present for the first time a sufficient criterion for the state distinction problem. Let $\|A\|_{\mathrm{F}}$ denote the Frobenius norm of a matrix A, i. e., $\|A\|_{\mathrm{F}} := \sqrt{\sum_{kl} |A_{kl}|^2}$. For a POVM $\mathcal{M}$ and quantum state σ in $\mathbb{C}^n$, let $\mathcal{M}(\sigma)$ denote the probability distribution on the outcomes of $\mathcal{M}$ got by measuring σ according to $\mathcal{M}$. Our main result can be stated informally as follows.

Result 1 (Informal statement). *Suppose σ_1, σ_2 are two quantum states in $\mathbb{C}^n$. Define $f := \|\sigma_1 - \sigma_2\|_{\mathrm{F}}$. If $\mathrm{rank}(\sigma_1) + \mathrm{rank}(\sigma_2)$ is not 'too large', then with probability at least $1 - \exp(-\Omega(\sqrt{n}) - \exp(-\Omega(f^2 n))$ over the choice of a random orthonormal basis $\mathcal{B}$ in $\mathbb{C}^n$, $\|\mathcal{B}(\sigma_1) - \mathcal{B}(\sigma_2)\|_1 > cf$, where c is a universal constant.*

Using the above result, we can show that if the minimum pairwise Frobenius distance of an ensemble $\mathcal{E} = \{\sigma_1, \ldots, \sigma_m\}$ of states in $\mathbb{C}^n$ is at least f, then with probability at least $1 - \exp(-n)$, a random POVM $\mathcal{F}$, with an appropriate notion of randomness, gives total variation distance at least cf between every pair of states of $\mathcal{E}$, where $c > 0$ is a universal constant. The notion of random POVM that we use is as follows: attach a zero ancilla in $\mathbb{C}^m$, where $m := \Theta\left(\frac{n \log^2 m}{f^2}\right)$, and measure $\sigma_i \otimes |0\rangle\langle 0|$ according to a random orthonormal basis in $\mathbb{C}^n \otimes \mathbb{C}^m$. In addition, as suggested by Result 1, if the maximum rank of a state in $\mathcal{E}$ is not too large, then we don't need a POVM at all, a random orthonormal basis in $\mathbb{C}^n$ will work just as well. We also construct examples of density matrices σ_1, σ_2 with $\|\sigma_1 - \sigma_2\|_{\mathrm{tr}} = 2$, where with very high probability the total variation distance given by a random POVM is at most $\sqrt{\|\sigma_1 - \sigma_2\|_{\mathrm{F}}}$, unless exponentially many ancilla qubits are used to define the random POVM.

Application to the HSP: Our random POVM method has information-theoretic implications about the HSP in a general group G. It is easy to see that the ensemble of coset states for subgroups of G is simultaneously block diagonal in the Fourier basis for G, where a block is labelled by an irreducible representation (irrep) of G and a row index. This leads us to consider the so-called *random Fourier method* for the HSP: apply the quantum Fourier transform over G to the given coset state and observe the name of an irrep ρ and a row index i, and then measure the resulting reduced state using a random POVM. Previously, a few examples of HSP's were given where random Fourier sampling required exponentially many copies of the coset state in order to identify the hidden subgroup with constant probability [GSVV04, MRRS04]. In these examples, the ranks of the blocks of the coset state in the Fourier basis were exponentially large. Using the fact that $\|A\|_{\mathrm{F}} \geq \frac{\|A\|_{\mathrm{tr}}}{\sqrt{\mathrm{rank}(A)}}$ for any matrix A, we prove a surprising positive counterpart to the above negative results. We show that polynomially many iterations of the random Fourier method give enough classical information to identify the hidden subgroup H if the ranks of the coset state in each block in the Fourier basis are polynomially bounded. In fact, we define a distance metric $r(H_1, H_2)$ between two subgroups $H_1, H_2 \leq G$ based on the Frobenius distance between the corresponding blocks of the coset states σ_{H_1} and σ_{H_2} in the Fourier basis of G, and show that random Fourier sampling gives total variation distance at least $\Omega(r(H_1, H_2))$ between σ_{H_1} and σ_{H_2} with exponentially high probability. If the ranks of the blocks of $\sigma_{H_1}, \sigma_{H_2}$ are polynomially bounded, then $r(H_1, H_2)$ is at least polynomially large. The previous work of [RRS05] also proposed a distance function $r'(H_1, H_2)$, but it was difficult to estimate $r'(H_1, H_2)$ except for very special cases. Also, the function $r'(H_1, H_2)$ is not powerful enough to even show that if the ranks of the blocks are $\sigma_{H_1}, \sigma_{H_2}$ are at most one, polynomially many iterations of random Fourier sampling suffice to identify the hidden subgroup with high probability. Our new result improves our understanding of the power of single register Fourier sampling, and establishes that the random POVM method can often be a powerful information-theoretic tool.

In particular, for the important special case when the hidden subgroup H forms a Gel'fand pair with the ambient group G, i. e., each block has rank either zero or one, $O(\log^3 |G|)$ iterations of random strong Fourier sampling give enough classical information to identify the hidden subgroup H with high probability. For many concrete examples e.g. affine group, Heisenberg group, the number of iterations of random Fourier sampling can be brought down to $O(\log |G|)$ by a more careful analysis. Gel'fand pairs have been studied extensively in group theory, and a lot of recent work [MR05] on the hidden subgroup problem has involved

Gel'fand pairs e.g. dihedral group [EH00, BCD05b] , affine group [MRRS04], Heisenberg group [RRS05, BCD05a]. For the dihedral and affine groups, it is possible to give explicit efficient measurement bases for the single register Fourier sampling procedure that identify the hidden subgroup with high probability using polynomially many copies. Interestingly, for the Heisenberg group no such explicit basis for single register Fourier sampling is known, though an explicit efficient entangled basis for *two-register* Fourier sampling is known [BCD05a]. The only proof that polynomially many iterations of single register Fourier sampling suffice information-theoretically to identify hidden subgroups in the Heisenberg group is through random Fourier sampling, and was first observed in [RRS05].

Since it can be shown that measuring in a Haar-random orthonormal basis is hard for a quantum computer, the main open question that arises from our work is whether there are efficiently implementable pseudo-random orthonormal bases for specific ensembles that have good distinguishing power. For example, such a basis for the representations of groups $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$, p prime, will give us algorithms for the HSP in those groups having an efficient quantum part followed by a possibly super polynomial classical post-processing. For super constant r, no such quantum algorithm is currently known. Current proposals of pseudo-random orthonormal bases [EWS$^+$03, ELL05] however, seem inadequate for our purposes.

Application to general state identification: Besides applications to the HSP, our random POVM method also has some interesting consequences for the general state identification problem. For an ensemble $\mathcal{E}$ of states in $\mathbb{C}^n$ with minimum pairwise trace distance δ and maximum rank r of a state, $t = O\left(\frac{r \log |\mathcal{E}|}{\delta^2}\right)$ independent copies of a state are enough to identify the state with high probability using t iterations of a random POVM. Since $r \leq n$, for a general ensemble of quantum states we get $t = O\left(\frac{n \log m}{\delta^2}\right)$ which is the first upper bound on the number of copies required for the general state identification problem to the best of our knowledge. For pure states, we get $t = O\left(\frac{\log m}{\delta^2}\right)$ which is optimal up to constant factors. This result for pure states can be independently proved by a detailed analysis of Gram-Schmidt orthonormalisation, but the resulting measurement is a *joint* measurement entangled across t registers. In contrast, note that all the state identification algorithms arising from our random POVM result are single register algorithms.

Related work: The so-called *pretty good measurement*, also known as the square-root measurement, has been proposed in the past as a measurement for the state identification problem [HW94]. Its performance is indeed 'pretty good' if the ensemble of states possesses some special symmetries; see e.g. [EMV04] and the references therein. The PGM approach has been recently applied to a few instances of the HSP also [BCD05b, BCD05a, MR05], showing that it maximises the probability of identifying the hidden subgroup for those instances. The PGM approach to state identification differs from our approach in an important way: the PGM approach does not usually give single register algorithms for state identification, whereas our approach based on state distinction does. This is because the PGM for t copies, in general, is a joint measurement and does not decompose as a tensor product of measurements on the individual copies. In fact, for the dihedral HSP studied in [BCD05b], an exponential number of iterations of the PGM for a single copy are required in order to identify a hidden reflection with constant probability. In contrast, polynomially many iterations of 'forgetful' Fourier sampling on single copies give enough classical information to identify a hidden reflection in the dihedral group [EH00].

Another problem similar to state distinction is as follows: for two a priori known ensembles $\mathcal{E}_1$, $\mathcal{E}_2$ of quantum states, is there a two-outcome POVM that identifies with reasonable probability to which ensemble a given state from $\mathcal{E}_1 \cup \mathcal{E}_2$ belongs? It turns out that the probability of error is related to the minimum trace distance between the convex hulls of $\mathcal{E}_1$ and $\mathcal{E}_2$ [GW05, Jai05], and is $1/2$ if the convex hulls intersect. In

contrast, in the state distinction problem we want to find a POVM with many outcomes that gives reasonable total variation distance between every pair of states of the ensemble. Having more than two outcomes allows us to find a pairwise distinguishing POVM even if the ensemble cannot be partitioned into two parts with disjoint convex hulls.

Proof technique: In order to show that, under suitable conditions, a random orthonormal basis $\mathcal{B}$ gives total variation distance at least $\Omega(\|\sigma_1 - \sigma_2\|_{\mathrm{F}})$ between two quantum states σ_1, σ_2, we have to analyse $\mathcal{B}$ in the eigenbasis of $\sigma_1 - \sigma_2$. Our techniques differ from earlier work on the power of random basis for state distinction [RRS05] in two different ways. First, the paper [RRS05] could not handle an arbitrary pair of quantum states σ_1, σ_2 because of using weaker symmetry arguments. Using better symmetry arguments and a new probabilistic analysis of the Gram-Schmidt orthonormalisation process, we overcome this limitation and reduce the problem to proving lower bounds on the tail of weighted sums of squares of Gaussian random variables. For the pairs of states considered in [RRS05], one only needed to prove tail lower bounds for an unweighted sum of squares of Gaussian, i. e., one needed to prove tail lower bounds for the chi-square distribution. The paper [RRS05] proved such bounds using the central limit theorem from probability theory. However, since we are now in the weighted case, the statement of the central limit theorem does not quite suffice. The main problem is that the central limit theorem cannot guarantee that a weighted sum of squared Gaussians exceeds its mean by a standard deviation with constant probability independent of the number of random variables and the weights. To do this, we have to use a powerful quantitative version of the central limit theorem known as the Berry-Esséen theorem combined with 'weight smoothening' arguments. This allows us to show that the tail of a weighted sum of squared Gaussian exceeds the ℓ_2-norm of the weight vector with constant probability. This is in contrast to Chernoff-like upper bounds on the tail of chi-square distributions that are more commonly seen in the study of measure concentration for random unitaries. Since the ℓ_2-norm of the weight vector is closely related to $\|\sigma_1 - \sigma_2\|_{\mathrm{F}}$, we get our main result easily after this. The Berry-Esséen theorem also indicates that a random orthonormal basis cannot achieve total variation distance much larger than $\|\sigma_1 - \sigma_2\|_{\mathrm{F}}$, and in fact, we give an example of states σ_1, σ_2 with trace distance 2 where a random basis cannot give total variation distance more than $\sqrt{\|\sigma_1 - \sigma_2\|_{\mathrm{F}}}$ with high probability.

2 Preliminaries

2.1 Measure concentration in $\mathbb{C}^n$

In this subsection, we prove some simple results about measure concentration phenomena in $\mathbb{C}^n$ for large n, that will be useful in the proof of our main theorem.

By a Gaussian probability distribution $\mathcal{G}$, we mean the one-dimensional real Gaussian probability distribution with mean 0 and variance 1, i. e., for $x \in \mathbb{R}$, the probability density of $\mathcal{G}$ at x is $\frac{e^{-x^2/2}}{\sqrt{2\pi}}$. We use $\Phi(\cdot)$ to denote the cumulative distribution function of $\mathcal{G}$, i .e., $\Phi(x)$ is the probability that $\mathcal{G}$ picks a real number less than or equal to x.

The following tail bound on the sum of squares of n independent Gaussians, also known as the chi-square distribution with n degrees of freedom, can be proved Chernoff-style using the moment generating function of the square of a Gaussian random variable.

Fact 1. *Let $G_1, \ldots, G_n$ be independent random variables where each G_i is distributed according to $\mathcal{G}$. Let $Y := \sum_{i=1}^n G_i^2$. For all $\epsilon \geq 0$,*

$$\Pr[Y > n(1+\epsilon)] < (\exp(-\epsilon/2) \cdot \sqrt{1+\epsilon})^n.$$

The same upper bound also holds for $\Pr[Y < n(1+\epsilon)]$ *when* $-1 < \epsilon < 0$.

Using Fact 1, we can prove the following lemma upper bounding the length of the projection of a random unit vector onto a fixed subspace.

Lemma 1. *Let* W *be a* k*-dimensional subspace of* $\mathbb{C}^n$*, where* $k \leq n/4$*. Let* v *be a random unit vector in* $\mathbb{C}^n$*. Let* Π_W *denote the orthonormal projector from* $\mathbb{C}^n$ *to* W*. Suppose* $4 \leq t \leq n/k$*. Then,*

$$\Pr\left[\|\Pi_W(v)\|^2 > t \cdot \frac{k}{n}\right] \leq \exp(-\Omega(tk)).$$

Also, for any $0 \leq \epsilon \leq 1/2$*,*

$$\Pr\left[(1-\epsilon)\frac{k}{n} \leq \|\Pi_W(v)\|^2 \leq (1+\epsilon)\frac{k}{n}\right] \geq 1 - \exp(-\Omega(\epsilon^2 k)).$$

Proof. We can choose a random unit vector $v \in \mathbb{C}^n$ as follows: choose a random vector $\hat{v} \in \mathbb{C}^n$ by choosing $2n$ independent real random variables $G_1, \ldots, G_{2n}$, where each G_i is distributed according to $\mathcal{G}$, and treating a complex number as a pair of real numbers. Now normalise $\hat{v}$ to get a random unit vector v; note that $\|\hat{v}\| = 0$ with probability 0. By symmetry, we can assume that W is spanned by the first k standard basis vectors in $\mathbb{C}^n$. Thus, $\|\Pi_W(v)\|^2 = \frac{\sum_{i=1}^{2k} G_i^2}{\sum_{j=1}^{2n} G_j^2}$. Using $\epsilon = -1/2$ in Fact 1, we get $\sum_{j=1}^{2n} G_j^2 > n$ with probability at least $1 - \exp(-\Omega(n))$ over the choice of v. Since $\exp(-\epsilon/2) \cdot \sqrt{1+\epsilon} \leq \exp(-\epsilon/10)$ for $\epsilon \geq 1$, using $\epsilon = t/4$ in Fact 1 we get $\sum_{i=1}^{2k} G_i^2 \leq \frac{(t+4)k}{2}$ with probability at least $1 - \exp(-\Omega(tk))$ over the choice of v. Thus, with probability at least $1 - \exp(-\Omega(tk)) - \exp(-\Omega(n))$ over the choice of v, $\|\Pi_W(v)\|^2 < \frac{(t+4)k}{2n} \leq \frac{tk}{n}$. This completes the proof of the first part of the lemma.

The proof of the second part of the lemma is very similar, using the inequality $\exp(-\epsilon/2) \cdot \sqrt{1+\epsilon} \leq -\epsilon^2/3$ for $0 \leq \epsilon \leq 1/2$. □

We now prove a lemma upper bounding the perturbation induced by the Gram-Schmidt orthonormalisation process on r random independent unit vectors in $\mathbb{C}^n$.

Lemma 2. *Let* $b'_1, \ldots, b'_r$ *be a sequence of random independent unit vectors in* $\mathbb{C}^n$*, where* $r \leq n$*. Let* $\tilde{b}_1, \ldots, \tilde{b}_r$ *be the corresponding sequence of unit vectors got by Gram-Schmidt orthonormalising* $b'_1, \ldots, b'_r$*. Fix* $M > 1$*. Then with probability at least* $1 - r \cdot \exp(-\Omega(Mr))$ *over the choice of* $b'_1, \ldots, b'_r$*,*

$$\||b'_i\rangle\langle b'_i| - |\tilde{b}_i\rangle\langle\tilde{b}_i|\|_{\mathrm{tr}} \leq O\left(\sqrt{\frac{Mr}{n}}\right)$$

for all $1 \leq i \leq r$*,*

Proof. For $1 \leq i \leq r$, let Π_i denote the orthonormal projector from $\mathbb{C}^n$ to the subspace spanned by $b'_1, \ldots, b'_i$. For $1 \leq i \leq r-1$, putting $t = \frac{Mr}{i}$ in the first part of Lemma 1, we see that with probability at least $1 - r\exp(-\Omega(Mr))$ over the choice of $b'_1, \ldots, b'_r$, $\|\Pi_i(b'_{i+1})\|^2 \leq O\left(\frac{Mr}{n}\right)$. Recall that $\tilde{b}_{i+1} := \frac{b'_{i+1} - \Pi_i(b'_{i+1})}{\|b'_{i+1} - \Pi_i(b'_{i+1})\|}$. Hence,

$$\begin{aligned}
\|\tilde{b}_{i+1} - b'_{i+1}\|^2 &= \|\Pi_i(b'_{i+1})\|^2 + \left(1 - \|b'_{i+1} - \Pi_i(b'_{i+1})\|\right)^2 \\
&= \|\Pi_i(b'_{i+1})\|^2 + \left(1 - \sqrt{1 - \|\Pi_i(b'_{i+1})\|^2}\right)^2 = 2 - 2\sqrt{1 - \|\Pi_i(b'_{i+1})\|^2} \\
&\leq 2 - 2\sqrt{1 - O\left(\frac{Mr}{n}\right)} \leq O\left(\frac{Mr}{n}\right).
\end{aligned}$$

The proposition now follows from the fact that for two unit vectors $|\psi\rangle$, $|\phi\rangle$, $\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_{\mathrm{tr}} \leq 2\||\psi\rangle - |\phi\rangle\|$. $\square$

We will require the following fact about the size of a δ-net in $\mathbb{C}^n$. A δ-net $\mathcal{N}$ is a finite set of unit vectors in $\mathbb{C}^n$ with the property that for any unit vector $v \in \mathbb{C}^n$, there exists a unit vector $v' \in \mathcal{N}$ such that $\|v - v'\| \leq \delta$. The fact follows from the proof technique of [Mat02, Lemma 13.1.1, Chapter 13] and by identifying $\mathbb{C}^n$ with $\mathbb{R}^{2n}$. Below for $1 \leq j \leq n$, e_j denotes the jth standard unit vector in $\mathbb{C}^n$, viz., the n-tuple containing a 1 in the jth location and zeroes elsewhere.

Fact 2. *Fix any $\delta \in (0, 1]$. Then, there is a δ-net $\mathcal{N}$ in $\mathbb{C}^n$ containing the n standard unit vectors $e_1, \ldots, e_n$ such that $|\mathcal{N}| \leq \left(\frac{4}{\delta}\right)^{2n}$.*

Using Fact 2, we can prove the following lemma upper bounding the spectral norm of an $n \times n$ matrix whose entries are independent random complex numbers with independent Gaussian real and imaginary parts.

Lemma 3. *Define a random $n \times n$ complex matrix M by independently choosing each entry to be a complex number whose real and imaginary parts are independently chosen according to the Gaussian distribution $\mathcal{G}$. Then, with probability at least $1 - \exp(-\Omega(n \log n))$ over the choice of M, $\|M\| \leq O(\sqrt{n \log n})$.*

Proof. Let $\delta := 1/\sqrt{n}$. Let $\mathcal{N}$ be a δ-net in $\mathbb{C}^n$ guaranteed by Fact 2. Fix any unit vector $v \in \mathbb{C}^n$. By symmetry, the probability distribution of $\|Mv\|^2$ is the same as that of $\|Me_1\|^2$, i. e., the probability distribution of $\|Mv\|^2$ is the same as that of the sum of squares of $2n$ independent Gaussians. Let $t := C \log n$, where C is a sufficiently large constant whose value will become clear later. Since $\exp(-\epsilon/2) \cdot \sqrt{1+\epsilon} \leq \exp(-\epsilon/10)$ for $\epsilon \geq 1$, using $\epsilon = t$ in Fact 1, we get that $\|Mv'\|^2 \leq (t+1)n$ for all $v' \in \mathcal{N}$ with probability at least $1 - (4\sqrt{n})^{2n} \cdot \exp(-\Omega(Cn \log n)) \geq 1 - \exp(-\Omega(n \log n))$ over the choice of M.

Note that for any vector $w \in \mathbb{C}^n$, we have

$$\begin{aligned}
\|Mw\|^2 &= \sum_{i=1}^{n} \left| \sum_{j=1}^{n} M_{ij} w_j \right|^2 \leq \sum_{i=1}^{n} \left(\sum_{j=1}^{n} |M_{ij}|^2 \right) \cdot \left(\sum_{j=1}^{n} |w_j|^2 \right) = \|w\|^2 \sum_{j=1}^{n} \sum_{j=1}^{n} |M_{ij}|^2 \\
&= \|w\|^2 \sum_{j=1}^{n} \|Me_j\|^2 \leq \|w\|^2 n^2 (t+1).
\end{aligned}$$

The inequality above follows from Cauchy-Schwartz. Now fix any unit vector $v \in \mathbb{C}^n$. Let v' be the closest vector to v from $\mathcal{N}$, where ties are broken arbitrarily. Thus, $\|v - v'\| \leq \delta$. We have

$$\begin{aligned}
\|Mv\|^2 &= \langle v|M^\dagger M|v\rangle = \langle v' + (v - v')|M^\dagger M|v' + (v - v')\rangle \\
&= \|Mv'\|^2 + \langle v'|M^\dagger M|v - v'\rangle + \langle v - v'|M^\dagger M|v'\rangle + \|M(v - v')\|^2 \\
&\leq \|Mv'\|^2 + 2\|Mv'\|\|M(v - v')\| + \|M(v - v')\|^2 \\
&\leq (t+1)n + 2\sqrt{(t+1)n} \cdot \|v - v'\| \cdot n\sqrt{t+1} + \|v - v'\|^2 n^2 (t+1) \\
&\leq (t+1)n + 2n^{3/2}(t+1)\delta + \delta^2 n^2 (t+1) \leq O(n \log n).
\end{aligned}$$

The first inequality above follows from Cauchy-Schwartz. The proof of the lemma is now complete. $\square$

Finally, we will require the following Berry-Esséen theorem from probability theory, which is a quantitative version of the central limit theorem [Fel71, Chapter XVI, Section 5, Theorem 2].

Fact 3 (Berry-Esséen theorem). *Let $X_1, \ldots, X_n$ be independent random variables. Define $\mu_i := \mathrm{E}[X_i]$, $\sigma_i := (\mathrm{E}[|X_i - \mu_i|^2])^{1/2}$, $\rho_i := (\mathrm{E}[|X_i - \mu_i|^3])^{1/3}$. Define the quantities*

$$\sigma^2 := \sum_{i=1}^n \sigma_i^2, \qquad \rho^3 := \sum_{i=1}^n \rho_i^3, \qquad X := \frac{1}{\sigma}\sum_{i=1}^n (X_i - \mu_i).$$

Then for all $x \in \mathbb{R}$,

$$|\Pr[X \le x] - \Phi(x)| \le \frac{6\rho^3}{\sigma^2}.$$

Remark: The constant 6 in the Berry-Esséen theorem can be improved; the current record is 0.7915 by Shiganov [Shi86]. However, Proposition 1 below holds as long as the constant is finite and independent of n and the random variables $X_1, \ldots, X_n$.

Using Fact 3, we prove the following proposition which will play a central role in the proof of our main theorem.

Proposition 1. *Let $G_1, \ldots, G_n$ be independent random variables where each G_i is distributed according to $\mathcal{G}$. Let $\lambda_1, \ldots, \lambda_n \in (0, 1]$. Define*

$$t := \sum_{i=1}^n \lambda_i, \qquad f := \sqrt{\sum_{i=1}^n \lambda_i^2}, \qquad X := \sum_{i=1}^n \lambda_i G_i^2.$$

Suppose $t \le 1$. Then, there is a constant c independent of n and $\lambda_1, \ldots, \lambda_n$ such that

$$\Pr[X > t + f] > c \qquad \text{and} \qquad \Pr[X < t] > c.$$

Proof. Without loss of generality, $\lambda_1 \ge \cdots \ge \lambda_n$. Let K_1 be a sufficiently large constant, whose choice will become clear later. Suppose $\lambda_1 \ge \frac{t}{K_1}$. Note that $\frac{t}{K_1} \le f \le t$. There is a constant c_1 depending on K_1 but independent of n and $\lambda_1, \ldots, \lambda_n$ such that $\Pr[G_1^2 > 2K_1] > c_1$, which implies that

$$\Pr[X > t + f] > \Pr[\lambda_1 G_1^2 > 2t] > \Pr\left[\frac{t}{K_1} G_1^2 > 2t\right] = \Pr\left[G_1^2 > 2K_1\right] > c_1.$$

Also,

$$\begin{aligned}
t = \mathrm{E}[X] &\ge t \cdot \Pr[t \le X \le t + f] + (t + f)\Pr[X > t + f] \\
&= t \cdot \Pr[X \ge t] + f \cdot \Pr[X > t + f] \\
&\ge t \cdot \Pr[X \ge t] + \frac{t}{K_1} \cdot c_1 \\
&= t \cdot (1 - \Pr[X < t]) + \frac{tc_1}{K_1} \\
\Rightarrow \quad & \Pr[X < t] \ge \frac{c_1}{K_1}.
\end{aligned}$$

Now, suppose $\lambda_1 < \frac{t}{K_1}$. Define independent random variables $X_i := \lambda_i G_i^2$. Let μ_i, σ_i, ρ_i be defined as in Fact 3. Recall that $\mathrm{E}[G_i^2] = 1$, $\mathrm{E}[|G_i^2 - 1|^2] = 2$ and that the absolute third central moment of G_i^2 is finite, say equal to K_2. Then,

$$\frac{6\sum_{i=1}^n \rho_i^3}{\sum_{i=1}^n \sigma_i^2} = \frac{6K_2 \sum_{i=1}^n \lambda_i^3}{2\sum_{i=1}^n \lambda_i^2} < \frac{6K_2 t}{2K_1} \le \frac{3K_2}{K_1}.$$

Taking $x = \frac{1}{\sqrt{2}}$ in Fact 3, we get

$$\Pr[X > t + f] \geq \left(1 - \Phi\left(\frac{1}{\sqrt{2}}\right)\right) - \frac{3K_2}{K_1}.$$

Similarly, taking $x = 0$ in Fact 3 we get

$$\Pr[X \leq t] \geq \Phi(0) - \frac{3K_2}{K_1} = \frac{1}{2} - \frac{3K_2}{K_1}.$$

Choosing K_1 to be a sufficiently large constant, we see that there exists a universal constant c_2 such that $\Pr[X > t + f] > c_2$ and $\Pr[X < t] = \Pr[X \leq t] > c_2$. Now letting $c := \min\left\{\frac{c_1}{K_1}, c_2\right\}$, we have that $\Pr[X > t + f] > c$ and $\Pr[X < t] > c$ always. Observe that c is a universal constant independent of n and $\lambda_1, \ldots, \lambda_n$. □

2.2 Quantum state distinction versus identification

In this subsection, we explore the connection between the problems of quantum state distinction and state identification.

A quantum state in $\mathbb{C}^n$ is modelled by a *density matrix* σ, which is an $n \times n$ Hermitian, positive semidefinite matrix with unit trace. A *positive operator-valued measure*, or POVM for short, is the most general measurement on quantum states. See e.g. [NC00] for a good introduction to density matrices and POVM's. A POVM $\mathcal{M}$ in $\mathbb{C}^n$ is a finite collection of positive operators E_i on $\mathbb{C}^n$, called elements of $\mathcal{M}$, that satisfy the completeness condition $\sum_i E_i = \mathbb{1}_n$. If the state of the quantum system is given by the density matrix σ, then the probability p_i to observe outcome labelled i is given by the Born rule $p_i = \mathrm{Tr}(\sigma E_i)$. We use $\mathcal{M}(\sigma)$ to denote the probability distribution on the outcomes of $\mathcal{M}$ got by measuring σ according to $\mathcal{M}$. The *trace norm* of an $n \times n$ matrix A is defined as $\|A\|_{\mathrm{tr}} := \mathrm{Tr}\sqrt{A^\dagger A}$. The *Frobenius norm* of A is defined as $\|A\|_{\mathrm{F}} := \sqrt{\mathrm{Tr} A^\dagger A}$, which is nothing but the ℓ_2-norm of the long vector in $\mathbb{C}^{n^2}$ corresponding to A. The following fact follows easily from the Cauchy-Schwartz inequality.

Fact 4. *For any matrix A, $\|A\|_{\mathrm{F}} \geq \frac{\|A\|_{\mathrm{tr}}}{\sqrt{\mathrm{rank}(A)}}$.*

Suppose there is an a priori known ensemble $\mathcal{E} = \{\sigma_1, \ldots, \sigma_m\}$ of quantum states in $\mathbb{C}^n$. Given t copies of a state σ_i, a *single register state identification* algorithm $\mathcal{A}$ for the ensemble $\mathcal{E}$ consists of a sequence of POVM's $\mathcal{F}_j$, $1 \leq j \leq t$, where $\mathcal{F}_j$ operates on the jth copy of σ_i. There is no bound on the number of outcomes of $\mathcal{F}_j$. The choice of $\mathcal{F}_j$ may depend on the observed outcomes of $\mathcal{F}_1, \ldots, \mathcal{F}_{j-1}$. After t observations, $\mathcal{A}$ does a classical post-processing and declares its guess for i. For all i, $1 \leq i \leq m$, we want $\mathcal{A}$ to guess i with probability at least $3/4$.

Let $0 \leq \delta \leq 2$. A POVM $\mathcal{M}$ for the *state distinction* problem with *distinguishing power* δ for the ensemble $\mathcal{E}$ is a POVM with the property that $\|\mathcal{M}(\sigma_i) - \mathcal{M}(\sigma_j)\|_1 \leq \delta$ for all $1 \leq i < j \leq m$. It is easy to see via the triangle inequality that if there exists a single register state identification POVM on t copies, then there exists a state distinction POVM with distinguishing power $\Omega(1/t)$. The following fact is a converse to the above observation; a proof sketch is included for completeness.

Fact 5. *Let $\mathcal{E} = \{\sigma_1, \ldots, \sigma_m\}$ be an a priori known ensemble of quantum states in $\mathbb{C}^n$. If there is a POVM $\mathcal{M}$ for the state distinction problem with distinguishing power δ for the ensemble $\mathcal{E}$, then there is a single register state identification algorithm $\mathcal{A}$ for ensemble $\mathcal{E}$ working on $t = O\left(\frac{\log m}{\delta^2}\right)$.*

Proof. Fix $1 \leq i < j \leq m$. Under the promise that the unknown state is either σ_i or σ_j, applying $\mathcal{M}$ to each of t copies of the unknown state followed by a maximum likelihood estimate identifies the correct state with probability at least $1 - \frac{1}{4m}$, as can be seen by a standard Chernoff bound. Let F_{ij} denote this maximum likelihood routine. The identification algorithm $\mathcal{A}$ starts by applying $\mathcal{M}$ on each of t copies of the unknown state, which a priori can be any $\sigma_i \in \mathcal{E}$. After that, $\mathcal{A}$ does $m-1$ iterations of a classical minimum-finding style post-processing procedure comparing two possible states σ_i, σ_j in an iteration, using the classical routines F_{ij} on the t observed outcomes. Note that the same t observed outcomes are reused by the various routines F_{ij}; no fresh measurements are done. The success probability of the minimum-finding style post-processing, and hence algorithm $\mathcal{A}$, is at least $1 - \frac{m-1}{4m} \geq 3/4$. □

2.3 Hidden subgroup problem and quantum Fourier transform

In this section, we explain the importance of the quantum Fourier transform as a means of attacking the hidden subgroup problem. For a general introduction to representation theory of finite groups, see e.g. [Ser77].

We use the term irrep to denote an irreducible unitary representation of a finite group G and denote by $\widehat{G}$ a complete set of inequivalent irreps. For any unitary representation ρ of G, let ρ^* denote the representation obtained by entry-wise conjugating the unitary matrices $\rho(g)$, where $g \in G$. Note that the definition of ρ^* depends upon the choice of the basis used to concretely describe the matrices $\rho(g)$. If ρ is an irrep of G so is ρ^*, but in general ρ^* may be inequivalent to ρ. Let V_ρ denote the vector space of ρ, define $d_\rho := \dim V_\rho$, and notice that $V_\rho = V_{\rho^*}$. The group elements $|g\rangle$, where $g \in G$ form an orthonormal basis of $\mathbb{C}^{|G|}$. Since $\sum_{\rho \in \widehat{G}} d_\rho^2 = |G|$, we can consider another orthonormal basis called the *Fourier basis* of $\mathbb{C}^{|G|}$ indexed by $|\rho, i, j\rangle$, where $\rho \in \widehat{G}$ and i, j run over the row and column indices of ρ. The quantum Fourier transform over G, QFT_G is the following linear transformation:

$$|g\rangle \mapsto \sum_{\rho \in \widehat{G}} \sqrt{\frac{d_\rho}{|G|}} \sum_{i,j=1}^{d_\rho} \rho_{ij}(g) |\rho, i, j\rangle.$$

It follows from Schur's orthogonality relations (see e.g. [Ser77, Chapter 2, Proposition 4, Corollary 3]) that QFT_G is a unitary transformation in $\mathbb{C}^{|G|}$.

For a subgroup $H \leq G$ and $\rho \in \widehat{G}$, define $\rho(H) := \frac{1}{|H|} \sum_{h \in H} \rho(h)$. It follows from Schur's lemma (see e.g. [Ser77, Chapter 2, Proposition 4]) that $\rho(H)$ is an orthogonal projection to the subspace of V_ρ consisting of vectors that are point-wise fixed by every $\rho(h)$, $h \in H$. Define $r_\rho(H) := \mathrm{rank}(\rho(H))$. Notice that $r_\rho(H) = r_{\rho^*}(H)$. The *standard method* of attacking the HSP in G using coset states [GSVV04] starts by forming the uniform superposition $\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle|0\rangle$. It then queries f to get the superposition $\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle|f(g)\rangle$. Ignoring the second register the reduced state on the first register becomes the density matrix $\sigma_H = \frac{1}{|G|} \sum_{g \in G} |gH\rangle\langle gH|$, that is the reduced state is a uniform mixture over all left coset states of H in G. It can be easily seen that applying QFT_G to σ_H gives us the density matrix $\frac{|H|}{|G|} \bigoplus_{\rho \in \widehat{G}} \bigoplus_{i=1}^{d_\rho} |\rho, i\rangle\langle\rho, i| \otimes \rho^*(H)$, where $\rho^*(H)$ operates on the space of column indices of ρ. Since the states σ_H are simultaneously block diagonal in the Fourier basis for any $H \leq G$, the elements of any POVM $\mathcal{M}$ operating on these states can without loss of generality be assumed to have the same block structure. From this it is clear that any distinguishing measurement without loss of generality first applies the quantum Fourier transform QFT_G to σ_H, measures the name ρ of an irrep, the index i of a row, and then measures the reduced state on the column space of ρ using a POVM $\mathcal{M}_\rho$ in $\mathbb{C}^{d_\rho}$. This POVM $\mathcal{M}_\rho$ may depend on ρ but is independent of i.

The probability of observing an irrep ρ in this quantum state is given by $\mathcal{P}_H(\rho) = \frac{d_\rho |H| r_\rho(H)}{|G|}$. Conditioned on observing ρ we obtain a uniform distribution $1/d_\rho$ on the row indices. The reduced state on the space of column indices after having observed an irrep ρ and a row index i is then given by the state $\rho^*(H)/r_\rho(H)$, and a basic task for a hidden subgroup finding algorithm is how to extract information about H from it. In this paper, we will investigate the case when $\mathcal{M}_\rho$ is a random POVM, for a suitable definition of randomness, in $\mathbb{C}^{d_\rho}$. We shall call this procedure *random Fourier sampling*. Grigni, Schulman, Vazirani and Vazirani [GSVV04] show that under certain conditions on G and H, random Fourier sampling gives exponentially small information about distinguishing H from the identity subgroup. In this paper, we prove a complementary information-theoretic result viz. under different conditions on G, $(\log |G|)^{O(1)}$ random strong Fourier samplings do give enough information to reconstruct the hidden subgroup H with high probability.

In *weak Fourier sampling*, we only measure the name of an irrep and ignore the reduced state on the column space. It can be shown [HRTS03] that for normal hidden subgroups H, no more information about H is contained in the reduced state. Thus, weak Fourier sampling is the optimal measurement to recover a normal hidden subgroup from its coset state. In particular, Fourier sampling is the optimal measurement on coset states for the abelian HSP.

Define a distance metric $w(H_1, H_2) := \|\mathcal{P}_{H_1} - \mathcal{P}_{H_2}\|_1 = \sum_{\rho \in \widehat{G}} |\mathcal{P}_{H_1}(\rho) - \mathcal{P}_{H_2}(\rho)|$ between subgroups $H_1, H_2 \leq G$. Adapting an argument in [HRTS03], it can be shown that $w(H_1, H_2) \geq 1/2$ if the *normal cores* of H_1 and H_2 are different [RRS05]. Recall that the normal core of a subgroup H is the largest normal subgroup of G contained in H. Thus, the main challenge is to distinguish between hidden subgroups H_1, H_2 from the same normal core family.

We next show that coset states corresponding to different hidden subgroups of a group have trace distance at least 1.

Proposition 2. *Let H_1, H_2 be different subgroups of a group G. Then, $\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{tr}} \geq 1$.*

Proof. For a subgroup $H \leq G$, we let G/H denote a complete set of left coset representatives of H in G. Since for any $c_1 \in G/H_1$,

$$|c_1 H_1\rangle = \sqrt{\frac{|H_1 \cap H_2|}{|H_1|}} \sum_{\substack{c \in G/(H_1 \cap H_2) \\ cH_1 = c_1 H_1}} |c(H_1 \cap H_2)\rangle,$$

we get

$$\sigma_{H_1} = \frac{|H_1|}{|G|} \sum_{c_1 \in G/H_1} |c_1 H_1\rangle\langle c_1 H_1| = \frac{|H_1 \cap H_2|}{|G|} \sum_{\substack{c, c' \in G/(H_1 \cap H_2) \\ cH_1 = c' H_1}} |c(H_1 \cap H_2)\rangle\langle c'(H_1 \cap H_2)|.$$

A similar fact is true for σ_{H_2}. We now define

$$\hat{\sigma}_{H_1} := \frac{|H_1 \cap H_2|}{|G|} \sum_{\substack{c, c' \in G/(H_1 \cap H_2) \\ cH_1 = c' H_1, c \neq c'}} |c(H_1 \cap H_2)\rangle\langle c'(H_1 \cap H_2)|.$$

We define $\hat{\sigma}_{H_2}$ similarly. Note that $\hat{\sigma}_{H_1}$, $\hat{\sigma}_{H_2}$ are Hermitian and for any $c \in G/(H_1 \cap H_2)$, $\langle c(H_1 \cap H_2)|\hat{\sigma}_{H_1}|c(H_1 \cap H_2)\rangle = 0$ and $\langle c(H_1 \cap H_2)|\hat{\sigma}_{H_1}|c(H_1 \cap H_2)\rangle = 0$.

We now observe that for any $c, c' \in G/(H_1 \cap H_2)$,

$$(\langle c(H_1 \cap H_2)|\sigma_{H_1}|c'(H_1 \cap H_2)\rangle \neq 0) \wedge (\langle c(H_1 \cap H_2)|\sigma_{H_2}|c'(H_1 \cap H_2)\rangle \neq 0) \iff c = c'.$$

This is because $cH_1 = c'H_1$ and $cH_2 = c'H_2$ implies that $c(H_1 \cap H_2) = c'(H_1 \cap H_2)$, i. e. $c = c'$. This implies that for any $c, c' \in G/(H_1 \cap H_2)$,

$$(\langle c(H_1 \cap H_2)|\hat{\sigma}_{H_1}|c'(H_1 \cap H_2)\rangle = 0) \vee (\langle c(H_1 \cap H_2)|\hat{\sigma}_{H_2}|c'(H_1 \cap H_2)\rangle = 0).$$

Thus, $\hat{\sigma}_{H_1}\hat{\sigma}_{H_2} = \hat{\sigma}_{H_2}\hat{\sigma}_{H_1} = 0$. Also, it follows that $\sigma_{H_1} - \sigma_{H_2} = \hat{\sigma}_{H_1} - \hat{\sigma}_{H_2}$.

Without loss of generality, H_1 is not a subgroup of H_2. Now,

$$\begin{aligned}
&\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{tr}} = \|\hat{\sigma}_{H_1} - \hat{\sigma}_{H_2}\|_{\mathrm{tr}} = \mathrm{Tr}\sqrt{(\hat{\sigma}_{H_1} - \hat{\sigma}_{H_2})^2} \\
&= \mathrm{Tr}\sqrt{\hat{\sigma}_{H_1}^2 + \hat{\sigma}_{H_2}^2 - \hat{\sigma}_{H_1}\hat{\sigma}_{H_2} - \hat{\sigma}_{H_2}\hat{\sigma}_{H_1}} \\
&= \mathrm{Tr}\sqrt{\hat{\sigma}_{H_1}^2 + \hat{\sigma}_{H_2}^2} \geq \mathrm{Tr}\sqrt{\hat{\sigma}_{H_1}^2} = \|\hat{\sigma}_{H_1}\|_{\mathrm{tr}}.
\end{aligned}$$

The inequality follows from the fact that $\hat{\sigma}_{H_1}^2$, $\hat{\sigma}_{H_2}^2$ are positive semidefinite operators and the square-root function is monotonically increasing for such operators. In order to evaluate $\|\hat{\sigma}_{H_1}\|_{\mathrm{tr}}$, notice that $\hat{\sigma}_{H_1} = \frac{|H_1 \cap H_2|}{|G|} \bigoplus_{c_1 \in G/H_1} M_{c_1}$, where for any $c_1 \in G/H_1$,

$$M_{c_1} := \sum_{\substack{c, c' \in G/(H_1 \cap H_2) \\ cH_1 = c'H_1 = c_1 H_1 \\ c \neq c'}} |c(H_1 \cap H_2)\rangle\langle c'(H_1 \cap H_2)|.$$

Now observe that M_{c_1} is of the form $J - I$, where J, I are the $\frac{|H_1|}{|H_1 \cap H_2|} \times \frac{|H_1|}{|H_1 \cap H_2|}$ all ones and identity matrices respectively. Hence, $\|M_{c_1}\|_{\mathrm{tr}} = 2\left(\frac{|H_1|}{|H_1 \cap H_2|} - 1\right)$ for all $c_1 \in G/H_1$. Thus,

$$\|\hat{\sigma}_{H_1}\|_{\mathrm{tr}} = \frac{|H_1 \cap H_2|}{|G|} \cdot \frac{|G|}{|H_1|} \cdot 2\left(\frac{|H_1|}{|H_1 \cap H_2|} - 1\right) = \frac{2(|H_1| - |H_1 \cap H_2|)}{|H_1|} \geq 1.$$

The inequality follows from the fact that $H_1 \cap H_2$ is a proper subgroup of H_1, since H_1 is not a subgroup of H_2. This completes the proof of the proposition. □

3 Random measurement bases and Frobenius distance

In this section, we prove our main result showing that a random POVM, for a suitable definition of randomness, distinguishes between two density matrices by at least their Frobenius distance with high probability. We first prove an important technical lemma that quickly implies our main theorem.

Lemma 4. *Let σ_1, σ_2 be two density matrices in $\mathbb{C}^n$. Define $f := \|\sigma_1 - \sigma_2\|_{\mathrm{F}}$. Then:*

1. *If $\mathrm{rank}(\sigma_1) + \mathrm{rank}(\sigma_2) \leq \sqrt{n}/K$, where K is a sufficiently large universal constant, then with probability at least $1 - \exp(-\Omega(\sqrt{n})) - \frac{\sqrt{n}}{K} \cdot \exp(-\Omega(f^2 n))$ over the choice of a random orthonormal measurement basis $\widehat{\mathcal{B}}$ in $\mathbb{C}^n$, $\|\widehat{\mathcal{B}}(\sigma_1) - \widehat{\mathcal{B}}(\sigma_2)\|_1 > \Omega(f)$;*

2. *Take a set $\mathcal{B}$ of n independent random vectors $\mathcal{B} := \{b_1, \ldots, b_n\}$ in $\mathbb{C}^n$, where each b_i is got by choosing n independent complex numbers whose real and imaginary parts are independently chosen according to the Gaussian $\mathcal{G}$. Define $\ell := \|\sum_{i=1}^n b_i b_i^\dagger\|$ and $\nu := \mathbb{1}_{\mathbb{C}^n} - \frac{1}{\ell}\sum_{i=1}^n b_i b_i^\dagger$. Let $\mathcal{M}$ denote the POVM on $\mathbb{C}^n$ consisting of the elements $\frac{b_i b_i^\dagger}{\ell}$ for $1 \leq i \leq n$, and the element ν. Note that $\mathcal{M}$ can be implemented as an orthonormal measurement in $\mathbb{C}^n \otimes \mathbb{C}^2$. Then with probability at least $1 - \exp(-\Omega(n))$ over the choice of $\mathcal{B}$, $\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 > \Omega\left(\frac{f}{\log n}\right)$.*

Proof. We start by proving the first part of the lemma. Define $t := \|\sigma_1 - \sigma_2\|_{\mathrm{tr}}$. We have $\mathrm{rank}(\sigma_1 - \sigma_2) \leq \sqrt{n}/K$, where K is a sufficiently large universal constant whose value will become clear later. Let $\widehat{\mathcal{B}} := \{|\hat{b}_1\rangle, \ldots, |\hat{b}_n\rangle\}$ be a random orthonormal basis of $\mathbb{C}^n$. Let $\widehat{\mathcal{B}}(\sigma_1)$, $\widehat{\mathcal{B}}(\sigma_2)$ denote the probability distributions on $[n]$ got by measuring σ_1, σ_2 respectively according to $\widehat{\mathcal{B}}$. Let $\lambda_1, \ldots, \lambda_k$ denote the positive eigenvalues, and $-\mu_{k+1}, \ldots, -\mu_{k+l}$ the negative eigenvalues of $\sigma_1 - \sigma_2$. Note that $k + l = \mathrm{rank}(\sigma_1 - \sigma_2) \leq \sqrt{n}/K$. We assume that we work in the eigenbasis of $\sigma_1 - \sigma_2$. Hence, we can write

$$\sigma_1 - \sigma_2 = \sum_{i=1}^{k} \lambda_i |i\rangle\langle i| - \sum_{j=k+1}^{k+l} \mu_j |j\rangle\langle j|, \qquad \sum_{i=1}^{k} \lambda_i = \sum_{j=k+1}^{k+l} \mu_j = \frac{t}{2}, \qquad \sum_{i=1}^{k} \lambda_i^2 + \sum_{j=k+1}^{k+l} \mu_j^2 = f^2.$$

Without loss of generality, $\sum_{i=1}^k \lambda_i^2 \geq \sum_{j=k+1}^{k+l} \mu_j^2 \Rightarrow \sum_{i=1}^k \lambda_i^2 \geq f^2/2$. Also, by the Cauchy-Schwartz inequality $t \leq f\sqrt{k+l}$. Then,

$$\begin{aligned}
\|\widehat{\mathcal{B}}(\sigma_1) - \widehat{\mathcal{B}}(\sigma_2)\|_1 &= \sum_{t=1}^{n} \left|\langle \hat{b}_t|\sigma_1|\hat{b}_t\rangle - \langle \hat{b}_t|\sigma_2|\hat{b}_t\rangle\right| = \sum_{t=1}^{n} \left|\langle \hat{b}_t|\sigma_1 - \sigma_2|\hat{b}_t\rangle\right| \\
&= \sum_{t=1}^{n} \left| \sum_{i=1}^{k} \lambda_i \left|\langle \hat{b}_t|i\rangle\right|^2 - \sum_{j=k+1}^{k+l} \mu_j \left|\langle \hat{b}_t|j\rangle\right|^2 \right|.
\end{aligned}$$

Define the random $n \times n$ unitary matrix $\widehat{\mathcal{B}}$ to be the matrix whose row vectors are $\langle \hat{b}_1|, \ldots, \langle \hat{b}_n|$. Then, $\|\widehat{\mathcal{B}}(\sigma_1) - \widehat{\mathcal{B}}(\sigma_2)\|_1 = \sum_{t=1}^n \left|\sum_{i=1}^k \lambda_i |\widehat{\mathcal{B}}_{ti}|^2 - \sum_{j=k+1}^{k+l} \mu_j |\widehat{\mathcal{B}}_{tj}|^2\right|$. Instead of generating the random unitary matrix $\widehat{\mathcal{B}}$ row-wise, we can generate it column-wise. The advantage now is that we only have to randomly generate the first $k+l$ orthonormal columns; the rest of the columns can be assumed to be zero without loss of generality. That is, we generate an $n \times (k+l)$ matrix $\widetilde{\mathcal{B}}$ whose columns are random orthonormal vectors $|\tilde{b}_1\rangle, \ldots, |\tilde{b}_{k+l}\rangle$ in $\mathbb{C}^n$. To generate the matrix $\widetilde{\mathcal{B}}$, we generate an $n \times (k+l)$ matrix $\mathcal{B}'$ whose columns are random independent unit vectors $|b'_1\rangle, \ldots, |b'_{k+l}\rangle$ in $\mathbb{C}^n$, and apply Gram-Schmidt orthonormalisation to get $|\tilde{b}_1\rangle, \ldots, |\tilde{b}_{k+l}\rangle$. Choosing $M = \frac{n}{K^2(k+l)^2}$ in Lemma 2, we get $\||\tilde{b}_t\rangle\langle\tilde{b}_t| - |b'_t\rangle\langle b'_t|\|_{\mathrm{tr}} < O\left(\frac{1}{K\sqrt{k+l}}\right)$ for all $1 \leq t \leq k+l$ with probability at least $1 - (k+l)\exp\left(-\Omega\left(\frac{n}{K^2(k+l)}\right)\right) \geq 1 - \exp(-\Omega(\sqrt{n}/K))$ over the choice of $\mathcal{B}'$. Let $\widetilde{\mathcal{B}}(\sigma_1) - \widetilde{\mathcal{B}}(\sigma_2)$ and $\mathcal{B}'(\sigma_1) - \mathcal{B}'(\sigma_2)$ denote the functions on $[n]$ defined by

$$(\widetilde{\mathcal{B}}(\sigma_1) - \widetilde{\mathcal{B}}(\sigma_2))(t) := \sum_{i=1}^k \lambda_i |\langle \tilde{b}_i|t\rangle|^2 - \sum_{j=k+1}^{k+l} \mu_j |\langle \tilde{b}_j|t\rangle|^2 = \sum_{i=1}^k \lambda_i |\widetilde{\mathcal{B}}_{ti}|^2 - \sum_{j=k+1}^{k+l} \mu_j |\widetilde{\mathcal{B}}_{tj}|^2,$$

$$(\mathcal{B}'(\sigma_1) - \mathcal{B}'(\sigma_2))(t) := \sum_{i=1}^k \lambda_i |\langle b'_i|t\rangle|^2 - \sum_{j=k+1}^{k+l} \mu_j |\langle b'_j|t\rangle|^2 = \sum_{i=1}^k \lambda_i |\mathcal{B}'_{ti}|^2 - \sum_{j=k+1}^{k+l} \mu_j |\mathcal{B}'_{tj}|^2$$

respectively, where $1 \leq t \leq n$. We now have

$$\|\widehat{\mathcal{B}}(\sigma_1) - \widehat{\mathcal{B}}(\sigma_2)\|_1 = \|\widetilde{\mathcal{B}}(\sigma_1) - \widetilde{\mathcal{B}}(\sigma_2)\|_1 = \sum_{t=1}^{n} \left| \sum_{i=1}^{k} \lambda_i |\langle \tilde{b}_i|t\rangle|^2 - \sum_{j=k+1}^{k+l} \mu_j |\langle \tilde{b}_j|t\rangle|^2 \right|$$

$$
\begin{aligned}
&\geq \sum_{t=1}^{n}\left|\sum_{i=1}^{k}\lambda_i|\langle b_i'|t\rangle|^2-\sum_{j=k+1}^{k+l}\mu_j|\langle b_j'|t\rangle|^2\right|-\sum_{t=1}^{n}\left|\sum_{i=1}^{k}\lambda_i\left(|\langle b_i'|t\rangle|^2-|\langle\tilde{b}_i|t\rangle|^2\right)\right|-\\
&\quad\sum_{t=1}^{n}\left|\sum_{j=k+1}^{k+l}\mu_j\left(|\langle b_j'|t\rangle|^2-|\langle\tilde{b}_j|t\rangle|^2\right)\right|\\
&\geq \|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1-\sum_{i=1}^{k}\lambda_i\sum_{t=1}^{n}\left||\langle b_i'|t\rangle|^2-|\langle\tilde{b}_i|t\rangle|^2\right|-\sum_{j=k+1}^{k+l}\mu_j\sum_{t=1}^{n}\left||\langle b_j'|t\rangle|^2-|\langle\tilde{b}_j|t\rangle|^2\right|\\
&\geq \|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1-\sum_{i=1}^{k}\lambda_i\||b_i'\rangle\langle b_i'|-|\tilde{b}_i\rangle\langle\tilde{b}_i|\|_{\mathrm{tr}}-\sum_{j=k+1}^{k+l}\mu_j\||b_j'\rangle\langle b_j'|-|\tilde{b}_j\rangle\langle\tilde{b}_j|\|_{\mathrm{tr}}\\
&\geq \|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1-O\left(\frac{1}{K\sqrt{k+l}}\right)\cdot\sum_{i=1}^{k}\lambda_i-O\left(\frac{1}{K\sqrt{k+l}}\right)\cdot\sum_{j=k+1}^{k+l}\mu_j\\
&= \|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1-t\cdot O\left(\frac{1}{K\sqrt{k+l}}\right)\geq\|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1-O\left(\frac{f}{K}\right)
\end{aligned}
$$

with probability at least $1-\exp\left(-\Omega\left(\frac{\sqrt{n}}{K}\right)\right)$ over the choice of $\mathcal{B}'$. The third inequality follows from the fact that the trace distance between two quantum states is an upper bound on the total variation distance between the probability distributions got by performing a measurement on the two states.

We generate $\mathcal{B}'$ by first generating an $n\times(k+l)$ matrix $\mathcal{B}$ whose entries are independent complex-valued random variables whose real and imaginary parts are each independently distributed according to the Gaussian $\mathcal{G}$, and then normalising each column of $\mathcal{B}$ in order to get $\mathcal{B}'$. Let $b_1,\ldots,b_{k+l}$ denote the columns of $\mathcal{B}$. Since $\exp(-\epsilon/2)\cdot\sqrt{1+\epsilon}\leq-\epsilon^2/3$ for $0\leq\epsilon\leq 1/2$, using $\epsilon=f/10$ in Fact 1 we see that with probability at least $1-(k+l)\exp(-\Omega(f^2n))$ over the choice of $\mathcal{B}$, $\|b_i\|^2\leq 2n\left(1+\frac{f}{10}\right)$ for $1\leq i\leq k$ and $\|b_j\|^2\geq 2n\left(1-\frac{f}{10}\right)$ for $k+1\leq j\leq k+l$. Consider any fixed t, $1\leq t\leq n$. By Proposition 1, with probability at least c^2 over the choice of $\mathcal{B}$,

$$
\sum_{i=1}^{k}\lambda_i|\mathcal{B}_{ti}|^2>2\sum_{i=1}^{k}\lambda_i+\sqrt{2\sum_{i=1}^{k}\lambda_i^2}\geq t+f\quad\text{and}\quad\sum_{j=k+1}^{k+l}\mu_j|\mathcal{B}_{ti}|^2<2\sum_{j=k+1}^{k+l}\mu_j=t.
$$

Call the above event E_t. If E_t occurs we have

$$
\begin{aligned}
&\left|\sum_{i=1}^{k}\lambda_i|\mathcal{B}_{ti}'|^2-\sum_{j=k+1}^{k+l}\mu_j|\mathcal{B}_{tj}'|^2\right|>\frac{t+f}{2n(1+\frac{f}{10})}-\frac{t}{2n(1-\frac{f}{10})}=-\frac{tf}{10n(1-\frac{f^2}{100})}+\frac{f}{2n(1+\frac{f}{10})}\\
&>\frac{f}{2n}\left(\frac{1}{1+\frac{\sqrt{2}}{10}}-\frac{2}{5}\right)>\frac{f}{6n}.
\end{aligned}
$$

Since the events E_t for different t are independent, using a standard Chernoff bound, with probability at least $1-\exp(-\Omega(n))$ over the choice of $\mathcal{B}$, at least $\frac{c^2n}{2}$ different t will satisfy the above inequality. This means that with probability at least $1-\exp(-\Omega(n))-(k+l)\exp(-\Omega(f^2n))$ over the choice of $\mathcal{B}$, $\|\mathcal{B}'(\sigma_1)-\mathcal{B}'(\sigma_2)\|_1\geq\frac{fc^2}{12}$. Thus, with probability at least $1-\exp(-\Omega(n))-(k+l)\exp(-\Omega(f^2n))-$

$\exp\left(-\Omega\left(\frac{\sqrt{n}}{K}\right)\right) \geq 1 - \exp\left(-\Omega\left(\frac{\sqrt{n}}{K}\right)\right) - \frac{\sqrt{n}}{K} \cdot \exp(-\Omega(f^2 n))$ over the choice of a random orthonormal basis $\widehat{\mathcal{B}}$ of $\mathbb{C}^n$, $\|\widehat{\mathcal{B}}(\sigma_1) - \widehat{\mathcal{B}}(\sigma_2)\|_1 > \frac{fc^2}{12} - O(f/K)$. Since c is a universal constant, we can choosing K to be a sufficiently large universal constant thus proving the first part of the lemma.

We now proceed to the proof of the second part of the lemma. Let $\lambda_1, \ldots, \lambda_k$ be the positive eigenvalues and $-\mu_{k+1}, \ldots, -\mu_n$ the non-positive eigenvalues of $\sigma_1 - \sigma_2$. By symmetry, we can assume that we are working in the eigenbasis of $\sigma_1 - \sigma_2$, i. e., the eigenbasis of $\sigma_1 - \sigma_2$ is the computational basis. Define the $n \times n$ matrix $\mathcal{B}$ to be the matrix whose column vectors are $b_1, \ldots, b_n$. Suppose v is a unit vector in $\mathbb{C}^m$. Then,

$$\langle v| \sum_{i=1}^{m} b_i b_i^\dagger |v\rangle = \sum_{i=1}^{n} |v^\dagger b_i|^2 = \|v^\dagger \mathcal{B}\|^2 = \|\mathcal{B}^\dagger v\|^2.$$

Hence we have

$$\ell = \|\sum_{i=1}^{n} b_i b_i^\dagger\| = \max_v \langle v| \sum_{i=1}^{m} b_i b_i^\dagger |v\rangle = \|\mathcal{B}^\dagger\|^2 = \|\mathcal{B}\|^2,$$

where the maximum is taken over all unit vectors $v \in \mathbb{C}^n$. The second equality follows because $\sum_{i=1}^{n} b_i b_i^\dagger$ is a positive matrix. By Lemma 3, $\ell = \|\mathcal{B}\|^2 \leq O(n \log n)$ with probability at least $1 - \exp(-\Omega(n \log n))$ over the choice of $\mathcal{B}$.

Now,

$$\begin{aligned}
\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 &= \frac{1}{\ell} \sum_{t=1}^{n} \left| b_t^\dagger \sigma_1 b_t - b_t^\dagger \sigma_2 b_t \right| + |\mathrm{Tr}(\sigma_1 \nu) - \mathrm{Tr}(\sigma_2 \nu)| \\
&\geq \frac{1}{\ell} \sum_{t=1}^{n} \left| b_t^\dagger \sigma_1 b_t - b_t^\dagger \sigma_2 b_t \right| = \frac{1}{\ell} \sum_{t=1}^{n} \left| \sum_{i=1}^{k} \lambda_i |b_t^\dagger |i\rangle|^2 - \sum_{j=k+1}^{n} \mu_j |b_t^\dagger |j\rangle|^2 \right| \\
&\geq \Omega\left(\frac{1}{n \log n}\right) \cdot \sum_{t=1}^{n} \left| \sum_{i=1}^{k} \lambda_i |\mathcal{B}_{it}|^2 - \sum_{j=k+1}^{n} \mu_j |\mathcal{B}_{jt}|^2 \right|.
\end{aligned}$$

By Proposition 1 and a standard Chernoff bound, we see that with probability at least $1 - \exp(-\Omega(n))$ over the choice of $\mathcal{B}$, for at least $\frac{c^2 n}{2}$ different t,

$$\left| \sum_{i=1}^{k} \lambda_i |\mathcal{B}_{it}|^2 - \sum_{j=k+1}^{n} \mu_j |\mathcal{B}_{jt}|^2 \right| > 2 \sum_{i=1}^{k} \lambda_i + \sqrt{2 \sum_{i=1}^{k} \lambda_i^2} - 2 \sum_{j=k+1}^{n} \mu_j \geq t + f - t = f.$$

Thus, with probability at least $1 - \exp(-\Omega(n)) - \exp(-\Omega(n \log n)) \geq 1 - \exp(-\Omega(n))$ over the choice of $\mathcal{B}$,

$$\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 > \Omega\left(\frac{1}{n \log n}\right) \cdot \frac{c^2 f}{n} = \Omega\left(\frac{f}{\log n}\right),$$

since c is a universal constant. Since the POVM $\mathcal{M}$ can be refined to a POVM with $2n$ rank one elements, $\mathcal{M}$ can be implemented as an orthonormal measurement in $\mathbb{C}^n \otimes \mathbb{C}^2$. This completes the proof of the second part of the lemma. □

We are now finally in a position to prove the main theorem of the paper.

Theorem 1. *Let σ_1, σ_2 be two density matrices in $\mathbb{C}^n$. Define $f := \|\sigma_1 - \sigma_2\|_{\mathrm{F}}$. Then:*

1. *Let $K > 1$ be a sufficiently large quantity. Consider an ancilla space $\mathbb{C}^m$ initialised to zero, where $m \geq \frac{4nK^2}{f^2}$. Let $\widehat{\mathcal{B}}$ be a random orthonormal measurement basis in $\mathbb{C}^n \otimes \mathbb{C}^m$. Let $\mathcal{M}$ denote the POVM on $\mathbb{C}^n$ got by attaching ancilla $|0\rangle$ to a state in $\mathbb{C}^n$ and applying the orthonormal measurement $\widehat{\mathcal{B}}$ in $\mathbb{C}^n \otimes \mathbb{C}^m$. Then with probability at least $1 - \exp(-\Omega(Kn))$ over the choice of $\widehat{\mathcal{B}}$, $\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 > \Omega(f)$;*

2. *Let $K \geq 1$ and define $m := Kn$. Take a set $\mathcal{B}$ of m independent random vectors $\mathcal{B} := \{b_1, \ldots, b_m\}$ in $\mathbb{C}^n \otimes \mathbb{C}^K$, where each b_i is got by choosing m independent complex numbers whose real and imaginary parts are independently chosen according to the Gaussian $\mathcal{G}$. Define $\ell := \|\sum_{i=1}^m b_i b_i^\dagger\|$ and $\nu := \mathbb{1}_{\mathbb{C}^n \otimes \mathbb{C}^K} - \frac{1}{\ell}\sum_{i=1}^m b_i b_i^\dagger$. Let $\mathcal{M}$ denote the POVM on $\mathbb{C}^n$ got by tensoring a zero ancilla over $\mathbb{C}^K$ to states in $\mathbb{C}^n$ and then performing the POVM $\overline{\mathcal{M}}$ in $\mathbb{C}^n \otimes \mathbb{C}^K$ consisting of the elements $\frac{b_i b_i^\dagger}{\ell}$ for $1 \leq i \leq m$, and the element ν. Note that $\mathcal{M}$ can be implemented as an orthonormal measurement in $\mathbb{C}^n \otimes \mathbb{C}^{2K}$. Then with probability at least $1 - \exp(-\Omega(m))$ over the choice of $\mathcal{B}$, $\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 > \Omega\left(\frac{f}{\log m}\right)$.*

Proof. In order to prove the first part of the theorem, let K be at least as large as the universal constant in the first part of Lemma 4. Thus, we start out with two density matrices $\bar{\sigma}_1 := \sigma_1 \otimes |0\rangle\langle 0|$, $\bar{\sigma}_2 := \sigma_2 \otimes |0\rangle\langle 0|$ in $\mathbb{C}^n \otimes \mathbb{C}^m$. Trivially, $\mathrm{rank}(\bar{\sigma}_1) + \mathrm{rank}(\bar{\sigma}_2) = \mathrm{rank}(\sigma_1) + \mathrm{rank}(\sigma_2) \leq 2n \leq \sqrt{nm}/K$. Also, $\|\bar{\sigma}_1 - \bar{\sigma}_2\|_{\mathrm{F}} = \|\sigma_1 - \sigma_2\|_{\mathrm{F}}$. By the first part of Lemma 4, with probability at least $1 - \exp(-\Omega(\sqrt{nm})) - \frac{\sqrt{nm}}{K} \cdot \exp(-\Omega(nmf^2)) \geq 1 - \exp(-\Omega(Kn))$ over the choice of a random orthonormal basis $\widehat{\mathcal{B}}$ of $\mathbb{C}^n \otimes \mathbb{C}^m$, $\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 = \|\widehat{\mathcal{B}}(\bar{\sigma}_1) - \widehat{\mathcal{B}}(\bar{\sigma}_2)\|_1 > \Omega(f)$. This completes the proof of the first part of the theorem.

A very similar strategy allows us to prove the second part of the theorem using the second part of Lemma 4. □

Remark: The point to note in the second part of the theorem is that the construction of the random POVM $\mathcal{M}$ does not require a priori knowledge of $\|\sigma_1 - \sigma_2\|_{\mathrm{F}}$. This will be useful in the application to the HSP, in the proof of Theorem 2

Finally, we present an example of a pair of density matrices σ_1, σ_2 where with high probability a random POVM cannot achieve a total variation distance much larger than $\sqrt{\|\sigma_1 - \sigma_2\|_{\mathrm{F}}}$, unless the dimension of the ancilla used by the POVM is exponentially larger than $\mathrm{rank}(\sigma_1) + \mathrm{rank}(\sigma_2)$. This is essentially because a sum of independent random variables cannot deviate from its mean by much more than its standard deviation.

Proposition 3. *Let σ_1, σ_2 be completely mixed states supported on two orthogonal r-dimensional subspaces of $\mathbb{C}^n$. Note that $\|\sigma_1 - \sigma_2\|_{\mathrm{F}} = \sqrt{2/r}$ and $\|\sigma_1 - \sigma_2\|_{\mathrm{tr}} = 2$. Let $\mathcal{B}$ be a random orthonormal basis in $\mathbb{C}^n$. Then, with probability at least $1 - n\exp(-\sqrt{r})$ over the choice of $\widehat{\mathcal{B}}$, $\|\mathcal{B}(\sigma_1) - \mathcal{B}(\sigma_2)\|_1 \leq O(r^{-1/4})$.*

Proof. Let $\mathcal{B} = \{|b_1\rangle, \ldots, |b_n\rangle\}$. Let W_1, W_2 denote the supports of σ_1, σ_2 respectively. Then, $\sigma_i = \frac{1}{r}\Pi_{W_i}$. Since each $|b_t\rangle$ is a random unit vector in $\mathbb{C}^n$, putting $\epsilon = Cr^{-1/4}$, C a universal constant whose value will become clear later, in the second part of Lemma 1, we get $\frac{1-\epsilon}{n} \leq \langle b_t|\sigma_i|b_t\rangle \leq \frac{1+\epsilon}{n}$ for $i = 1, 2$ and all $1 \leq t \leq n$, with probability at least $1 - n\exp(-\sqrt{r})$ over the choice of $\mathcal{B}$. Thus,

$$\|\mathcal{B}(\sigma_1) - \mathcal{B}(\sigma_2)\|_1 = \sum_{t=1}^n |\langle b_t|\sigma_1|b_t\rangle - \langle b_t|\sigma_2|b_t\rangle| \leq \sum_{t=1}^n \frac{2\epsilon}{n} \leq 2\epsilon.$$

This completes the proof of the proposition. $\square$

Now, if we think of $\mathbb{C}^n$ as $\mathbb{C}^{2r} \otimes \mathbb{C}^m$, where $m := \frac{n}{2r}$, we see that a random POVM in $\mathbb{C}^{2r}$ cannot distinguish between σ_1, σ_2 by much more than $\sqrt{\|\sigma_1 - \sigma_2\|_{\mathrm{F}}}$, unless n is exponentially large compared to $\mathrm{rank}(\sigma_1) + \mathrm{rank}(\sigma_2)$.

4 Random measurement bases and the HSP

In this section, we study the implications of Theorem 1 for the hidden subgroup problem.

Theorem 1 is in most cases not immediately useful in obtaining single register algorithms for the HSP. This is because for two candidate hidden subgroups H_1, H_2, $\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{F}} \leq \|\sigma_{H_1}\|_{\mathrm{F}} + \|\sigma_{H_2}\|_{\mathrm{F}} = \sqrt{\frac{|H_1|}{|G|}} + \sqrt{\frac{|H_2|}{|G|}}$. Thus, even though $\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{tr}} \geq 1$ by Proposition 2, $\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{F}}$ can be exponentially small if $|H_1|$, $|H_2|$ are exponentially small compared to $|G|$. In most examples of interest this is indeed the case. Fortunately, we can make good use of the fact that the coset states for different subgroups of G are simultaneously block diagonal in the Fourier basis of G. Hence, we investigate the power of random Fourier sampling in distinguishing between coset states. The advantage of this is that after doing the quantum Fourier transform and measuring an irrep name and a row index, we may be left with a reduced state on the space of column indices with polynomially bounded rank. If this happens, the average Frobenius distance between the blocks of σ_{H_1} and σ_{H_2} will be polynomially large even though $\|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{F}}$ may be exponentially small. In fact, for several cases of the HSP studied in the literature, the rank of the reduced state is in fact either 0 or 1 i. e., the hidden subgroup forms a Gel'fand pair with the ambient group.

To make the above reasoning precise, we define a new distance metric between two coset states σ_{H_1}, σ_{H_2}. Below, we use the notation of Section 2.3.

Definition 1 ($r(H_1, H_2)$). *Let G be a group and $H_1, H_2 \leq G$. Define*

$$r(H_1, H_2) := w(H_1, H_2) + \frac{1}{|G| \log |G|} \cdot \sum_{\rho \in \widehat{G}} d_\rho \, \||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{F}}$$

The importance of $r(H_1, H_2)$ follows from the following theorem.

Theorem 2. *Let G be a group and $H_1, H_2 \leq G$. Let $\mathcal{M}$ denote the POVM corresponding to the following random Fourier sampling procedure: apply QFT_G to the given coset state, measure the name of an irrep $\rho \in \widehat{G}$ and a row index i, and then apply a random POVM $\mathcal{M}_\rho$ on the resulting reduced state on the space of column indices, where $\mathcal{M}_\rho$ is defined as in the second part of Theorem 1 with $K_\rho := \left\lceil \frac{C \log^2 |G|}{d_\rho} \right\rceil$, where C is a sufficiently large universal constant. Then with probability at least $1 - \exp(-\log^2 |G|)$ over the choice of $\mathcal{M}$, $\|\mathcal{M}(\sigma_{H_1}) - \mathcal{M}(\sigma_{H_2})\|_1 \geq \Omega(r(H_1, H_2))$.*

Proof. Let σ_1, σ_2 be two quantum states and $p_1, p_2 \geq 0$. Suppose $p_1 \geq p_2$. Then,

$$\|p_1\sigma_1 - p_2\sigma_2\|_{\mathrm{F}} \leq \|p_1(\sigma_1 - \sigma_2)\|_{\mathrm{F}} + \|(p_1 - p_2)\sigma_2\|_{\mathrm{F}} \leq p_1\|\sigma_1 - \sigma_2\|_{\mathrm{F}} + |p_1 - p_2|.$$

Now,

$$\|p_1\mathcal{M}(\sigma_1) - p_2\mathcal{M}(\sigma_2)\|_1 = \|p_1(\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)) + (p_1 - p_2)\mathcal{M}(\sigma_2)\|_1 \geq \frac{p_1}{2}\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1.$$

The inequality above follows by considering those outcomes of $\mathcal{M}$ that have at least as much probability for σ_1 as for σ_2, and the fact that $(p_1 - p_2)\mathcal{M}(\sigma_2)$ is a vector with non-negative entries. Also, $\|p_1\mathcal{M}(\sigma_1) - p_2\mathcal{M}(\sigma_2)\|_1 \geq |p_1 - p_2|$. Now suppose $\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 \geq \frac{\|\sigma_1 - \sigma_2\|_{\mathrm{F}}}{L}$, where $L \geq 1$. Then,

$$\begin{aligned}
\|p_1\mathcal{M}(\sigma_1) - p_2\mathcal{M}(\sigma_2)\|_1 &\geq \frac{|p_1 - p_2|}{2} + \frac{p_1}{4}\|\mathcal{M}(\sigma_1) - \mathcal{M}(\sigma_2)\|_1 \\
&\geq \frac{|p_1 - p_2|}{4L} + \frac{p_1}{4L}\|\sigma_1 - \sigma_2\|_{\mathrm{F}} \geq \frac{\|p_1\sigma_1 - p_2\sigma_2\|_{\mathrm{F}}}{4L}.
\end{aligned}$$

Now suppose we apply QFT_G and measure an irrep name ρ and a row index i. We apply the above reasoning to the random POVM M_ρ with $L = \log|G|$. Using the second part of Theorem 1, we get that with probability at least $1 - \exp(-\log^2|G|)$ over the choice of M_ρ, $\|\mathcal{M}_\rho(\rho(H_1)) - \mathcal{M}_\rho(\rho(H_2))\|_1 \geq \Omega\left(\frac{\||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{F}}}{\log|G|}\right)$. Hence for the random Fourier sampling POVM $\mathcal{M}$, with probability at least $1 - \exp(-\log^2|G|)$ over the choice of $\mathcal{M}$,

$$\|\mathcal{M}(\sigma_{H_1}) - \mathcal{M}(\sigma_{H_2})\|_1 \geq \Omega\left(\frac{1}{|G|\log|G|} \cdot \sum_{\rho\in\widehat{G}} d_\rho \, \||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{F}}\right).$$

The theorem now follows because random Fourier sampling always does at least as well as weak Fourier sampling. □

The following corollary is now easy to prove.

Corollary 1. *Let G be a group. Suppose for every irrep $\rho \in \widehat{G}$ and subgroup $H \leq G$, $\mathrm{rank}(\rho(H)) \leq (\log|G|)^{O(1)}$. Then the random Fourier method of Theorem 2 gives rise to a single register algorithm identifying with probability at least $3/4$ the hidden subgroup H from $(\log|G|)^{O(1)}$ copies of σ_H.*

Proof. Consider two distinct subgroups $H_1, H_2 \leq G$. Since coset states are block diagonal in the Fourier basis of G, using Theorem 2, Proposition 2 and Fact 4 we get

$$\begin{aligned}
1 \leq \|\sigma_{H_1} - \sigma_{H_2}\|_{\mathrm{tr}} &= \sum_{\rho\in\widehat{G}} \frac{d_\rho}{|G|}\||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{tr}} \\
&\leq \sum_{\rho\in\widehat{G}} \frac{d_\rho}{|G|}\||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{F}} \cdot (\mathrm{rank}(\rho(H_1)) + \mathrm{rank}(\rho(H_2))) \\
&\leq (\log|G|)^{O(1)} \cdot \left(\sum_{\rho\in\widehat{G}} \frac{d_\rho}{|G|}\||H_1|\rho(H_1) - |H_2|\rho(H_2)\|_{\mathrm{F}}\right) \\
&\leq (\log|G|)^{O(1)} \cdot r(H_1, H_2).
\end{aligned}$$

Let $\mathcal{M}$ denote the random Fourier sampling POVM of Theorem 2. Then with probability at least $1 - \exp(-\log^2|G|)$ over the choice of $\mathcal{M}$, $\|\mathcal{M}(\sigma_{H_1}) - \mathcal{M}(\sigma_{H_2})\|_1 \geq \Omega(r(H_1, H_2)) \geq (\log|G|)^{-O(1)}$. Since a group G can have at most $2^{\log^2|G|}$ subgroups, by the union bound on probabilities, with probability at least $1 - \exp(-\Omega(\log^2|G|))$ over the choice of $\mathcal{M}$, $\|\mathcal{M}(\sigma_{H_1}) - \mathcal{M}(\sigma_{H_2})\|_1 \geq (\log|G|)^{-O(1)}$ for all subgroups $H_1, H_2 \leq G$. The corollary now follows from Fact 5. □

Finally, we remark that in many important examples of the HSP where most of the probability lies on high dimensional irreps and the blocks corresponding to these irreps have low rank, one can save a factor of $\log |G|$ in the denominator of the definition of $r(H_1, H_2)$ and prove Theorem 2 with this improved definition of $r(H_1, H_2)$. This improvement follows by using the first part of Lemma 4 instead of the second part of Theorem 1 in the proof of Theorem 2. Such a saving can be done, for example, for suitable subgroups of the affine group, Heisenberg group and groups $\mathbb{Z}_p^r \rtimes \mathbb{Z}_p$, p prime, $r \geq 2$.

5 The general state identification problem

In this section, we study the implications of Theorem 1 to the state identification problem for a general ensemble of quantum states. To the best of our knowledge, this problem does not seem to have been studied before. The following theorem with $r = n$ gives an upper bound on the number of copies required to identify a given state information-theoretically with high probability for any ensemble.

Theorem 3. *Let $\mathcal{E} = \{\sigma_1, \ldots, \sigma_k\}$ be an a priori known ensemble of quantum states in $\mathbb{C}^n$. Suppose the minimum trace distance between a pair of states from $\mathcal{E}$ is at least t. Let r denote the maximum rank of a state in $\mathcal{E}$. Then, there is a POVM $\mathcal{M}$ in $\mathbb{C}^n$ such that $\mathcal{M}^{\otimes \ell}$ acting on $\sigma_i^{\otimes \ell}$ gives enough classical information to identify i with probability at least $3/4$, where $\ell = O\left(\frac{r \log k}{t^2}\right)$.*

Proof. Define $f := \frac{t}{\sqrt{r}}$. Let $\mathcal{M}$ be the random POVM guaranteed by the second part of Theorem 1 with $m := \frac{16nK^2 \log^2 m}{f^2}$. Fix any pair of states σ_i, σ_j, $i \neq j$ from $\mathcal{E}$. Then with probability at least $1 - \exp(-\Omega(8n \log m)) \geq 1 - \frac{1}{m^2}$ over the choice of $\mathcal{M}$,

$$\|\mathcal{M}(\sigma_i) - \mathcal{M}(\sigma_j)\|_1 > \Omega(\|\sigma_i - \sigma_j\|_{\mathrm{F}}) \geq \Omega\left(\frac{\|\sigma_i - \sigma_j\|_{\mathrm{tr}}}{\sqrt{\mathrm{rank}(\sigma_i - \sigma_j)}}\right) \geq \Omega\left(\frac{t}{\sqrt{r}}\right).$$

By the union bound on probabilities, there is a POVM $\mathcal{M}$ on $\mathbb{C}^n$ such that the above inequality holds for every pair of states from $\mathcal{E}$. By Fact 5, applying $\mathcal{M}^{\otimes \ell}$ on $\sigma_i^{\otimes l}$, where $\ell = O\left(\frac{r \log k}{t^2}\right)$ gives enough classical information to identify i with probability at least $3/4$. □

Acknowledgements

The author wishes to thank Andris Ambainis, Martin Rötteler, Debbie Leung, Joseph Emerson and Christoph Dankert for useful discussions, and Jaikumar Radhakrishnan for feedback on an earlier version of the paper.

References

[BCD05a] D. Bacon, A. Childs, and W. van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, 2005. Also arxiv preprint quant-ph/0504083.

[BCD05b] D. Bacon, A. Childs, and W. van Dam. Optimal measurements for the dihedral hidden subgroup problem. Arxiv preprint quant-ph/0501044, 2005.

[EH00] M. Ettinger and P. Høyer. On quantum algorithms for noncommutative hidden subgroups. *Advances in Applied Mathematics*, 25(3):239–251, 2000.

[ELL05] J. Emerson, E. Livine, and S. Lloyd. Convergence conditions for random quantum circuits. Arxiv preprint quant-ph/0503210, 2005.

[EMV04] Y. Eldar, A. Megretski, and G. Verghese. Optimal detection of symmetric mixed quantum states. *IEEE Transactions on Information Theory*, 50:1198–1207, 2004.

[EWS+03] J. Emerson, Y. Weinstein, M. Saraceno, S. Lloyd, and D. Cory. Pseudo-random unitary operators for quantum information processing. *Science*, 302:2098–2100, 2003.

[Fel71] W. Feller. *Probability Theory and its Applications*, volume 2. John Wiley and Sons, 1971.

[GSVV04] M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica*, pages 137–154, 2004.

[GW05] G. Gutoski and J. Watrous. Quantum interactive proofs with competing provers. In *Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, vol. 3404, pages 605–616. Springer-Verlag, 2005. Also arxiv preprint cs.CC/0412102.

[HRTS03] S. Hallgren, A. Russell, and A. Ta-Shma. The hidden subgroup problem and quantum computation using group representations. *SIAM Journal on Computing*, 32(4):916–934, 2003.

[HW94] P. Hausladen and W. Wootters. A 'pretty good' measurement for distinguishing quantum states. *Journal of Modern Optics*, 41:2385–2390, 1994.

[Jai05] R. Jain. Distinguishing sets of quantum states. Arxiv preprint quant-ph/0506205, 2005.

[Mat02] J. Matoušek. *Lectures on Discrete Geometry*. Graduate Texts in Mathematics. Springer-Verlag, 2002.

[MR05] C. Moore and A. Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. Arxiv preprint quant-ph/0501177, 2005.

[MRRS04] C. Moore, D. Rockmore, A. Russell, and L. Schulman. The power of basis selection in Fourier sampling: Hidden subgroup problems in affine groups. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1113–1122, 2004. Journal version in preparation. Also arxiv preprint quant-ph/0503095.

[MRS05] C. Moore, A. Russell, and L. Schulman. The symmetric group defies strong Fourier sampling. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 479–488, 2005. Also arxiv preprint quant-ph/0501056.

[NC00] M. Nielsen and I. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.

[RRS05] J. Radhakrishnan, M. Rötteler, and P. Sen. On the power of random bases in Fourier sampling: Hidden subgroup problem in the Heisenberg group. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming*, Lecture Notes in Computer Science, vol. 3580, pages 1399–1411. Springer-Verlag, 2005. Also arxiv preprint quant-ph/0503114.

[Ser77] J-P. Serre. *Linear representations of finite groups*. Graduate Texts in Mathematics. Springer-Verlag, 1977.

[Shi86] I. Shiganov. Refinement of the upper bound of the constant in the central limit theorem. *Journal of Soviet Mathematics*, pages 2545–2550, 1986.

[Sho97] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.